

-1-

| | |
|----------------------|--|
| Date: <u>8-28-00</u> | Express Mail Label No. <u>EL387775583 US</u> |
|----------------------|--|

Inventor(s): Chandan Adhikari, John Anderson, Mark E. Bucci, Paul J.
Piccolomini, Dennis Smith and Joseph A. Velazquez
Attorney's Docket No.: 1768.2001-001

BUSINESS VOLUME AND WORKFORCE REQUIREMENTS FORECASTER

RELATED APPLICATIONS

5 This application claims benefit to U.S. Provisional Application Number
60/_____, filed June 28, 2000, entitled "Business Volume and Workforce
Requirements Forecaster Engine," by Chandan Adhikari, John Anderson, Mark E.
Bucci, Paul J. Piccolomini, Dennis Smith and Joseph A. Velazquez, the entire teachings
of which are incorporated herein by reference.

BACKGROUND OF THE INVENTION

10 The scheduling of employee workshifts is typically based on work requirements,
and is usually done in incremental time periods such as 15-minute periods.

For example, a supermarket may need four baggers and five cashiers between
9:15 AM and 1:15 PM, and only two baggers and three cashiers from 1:15 PM until
5:30 PM, at which time the need may again be elevated. Of course, scheduling must be
15 satisfied from the existing pool of employees. At the same time, many rules or
constraints must also be satisfied. These rules include minimum and maximum hours
for a given employee, rules related to employment of minors, rules relating to break
times, rules dictated by union contracts, etc.

Workforce requirements, that is, the specified need for so many workers to fill a
20 certain time slot for a specific job or task, are typically provided either by a manager's
best guess, or by some forecasting mechanism. Forecasters typically look at historical

009330-22164960

data from corresponding or similar periods from the past to produce an estimate of expected workload requirements. These requirements can then be fed to the scheduler.

SUMMARY OF THE INVENTION

A complete workforce management system allows users to forecast, schedule,
5 track, and pay their workforce according to their business' expected sales mix, balancing the need for maintaining strong customer satisfaction with the complicated scheduling and payroll demands of the workforce.

This is accomplished through the implementation of three modules: a forecaster, a dynamic scheduler, and a time and attendance module.

10 The forecaster, the invention described herein, processes and analyses a site's historical business mix, and provides an accurate volume and labor, i.e., workforce requirements, forecasts based on that data.

The dynamic scheduler analyzes workforce requirements in conjunction with scheduling rules required by the workplace, as well as the constraints of individual
15 employees. It generates a schedule which best considers these three factors according to user preferences. The dynamic scheduler is described in U.S. Application S/N 09/599,016, filed June 21, 2000, which is incorporated by reference herein in its entirety.

The Time and Attendance module tracks time and pay by employee, and
20 calculates pay cards per employee based on the labor rules set by the user. A Time and Attendance module is described in U.S. Application S/N 09/524,310, filed March 14, 2000, which is incorporated by reference herein in its entirety.

The forecaster determines workforce requirements by job by period of the day in such a way as to schedule an optimal number of people. It does this by using historical
25 business volume data to create a volume forecast, distributing the volume forecast over periods of the day, and applying labor standards by job to the distributed volume to

009230-22764960

determine workforce requirements. The forecaster is supported by an extensive reporting system and comparative data feedback from the Time and Attendance module.

A “volume” is a trackable value of business activity at the user’s site, such as sales or number of items sold. The forecaster tracks the actual values, i.e., the volumes, that have occurred over time, and uses that history when it creates a volume forecast. It calculates a labor forecast based on the volume forecast.

In accordance with the present invention, a method for forecasting business volume and workforce requirements includes defining both a business structure and a forecast structure, wherein certain hierarchical levels of the forecast structure map to corresponding hierarchical levels in the business structure. Business volume is forecast for a predefined duration, responsive to both a first set of historical data, and to the business and forecast structures.

A first portion of the data from the first set of historical data may be archived on a per day basis, while a second portion of the data is archived on a per period basis, where a period may cover, for example, fifteen minutes.

Forecasting business volume can be done using a daily trend forecasting algorithm, or using an exponential smoothing algorithm. One algorithm may be used for some forecast categories while the other algorithm is used for other categories.

In at least one embodiment, forecasting business volume comprises forecasting
20 daily quantities over a predefined duration.

In addition, forecasting business volume can be performed at plural levels of the forecast structure.

In at least one embodiment, hierarchical levels of the forecast structure which map to corresponding hierarchical levels in the business structure are location, department and job. A location can be subdivided into a plurality of sub-locations.

Certain hierarchical levels in the forecast structure can be at different depths within the forecast structure than the corresponding hierarchical levels in the business structure.

At least one embodiment of the present invention includes forecasting a traffic pattern for the predefined duration, responsive to a second set of historical data. The second set of historical data can be independent of the first set, or can partially or wholly overlap with the first set.

5 The forecast traffic pattern is preferably a composite of historical data from a plurality of selected dates, which are selected by finding a predetermined number of dates which best match designated criteria. Such criteria, which can be weighted, include same day of week, nearest day, event ratio and same open/close time.

Preferably, workforce requirements are calculated for the predefined duration,
10 based on the forecast business volume and on the forecast traffic pattern.

Calculation of workforce requirements can include task level consolidation. Specific tasks within a job are scheduled, where each task is associated with a standard, and consolidating the scheduled tasks are consolidated into a job schedule. The decision to apply a standard can be based, for example, on economy of scale.

15 In at least one embodiment, calculating workforce requirements includes resource leveling. Resource leveling includes repeatedly determining valleys in a preliminary schedule, ranking the valleys, assigning at least one unassigned task to a highest-ranked valley. Valleys can be ranked, for example, based on their depth and width. For example, a valley's rank may be computed as $(D/W) * C$, where D is the
20 valley's depth, W is the valley's width, and C is the valley's rounding cost. Valleys can be determined by first determining peaks in the preliminary schedule. Preferably, at least one unassigned task is assigned to a lowest portion of the highest-ranked valley.

The calculation of workforce requirements can include dynamic standard assignment, wherein different metrics are selected at different times, and where a task is associated with a plurality of standards. The selection of metrics at a specific time can be responsive to conditions, for example, weather conditions such as outdoor temperature, at the specific time.

In at least one embodiment, an event calendar is defined, and at least one event is selected from the event calendar such that the event is considered in the step of forecasting. For example, if the selected event does not occur during the forecast period, its influence is removed from the forecast if the event occurred during a
5 corresponding period from which the historical data was taken. On the other hand, if the selected event does occur during the forecast period, its influence is added to the forecast if the event did not occur during a corresponding period from which the historical data was taken, since such influence would by default be missing in the forecast due to the non-occurrence of the event in the historical data.

- 10 An event can be associated with at least one category in the forecast structure, and a plurality of events can be selected for a particular day.

When calculating forecast values for an upcoming day marked with some event, a search is conducted for historical dates marked with the same event marker. Upon finding such a date, a ratio is calculated of volume activity associated with the historical
15 date to the volume activity of plural days surrounding the historical date. A forecast for the upcoming forecast day is calculated as if the forecast day were a normal, non-event day. Finally, the forecast is adjusted by the calculated ratio. This can be done for each business volume.

- 20 Preferably, business volume types are user-definable, and can include any or all of sales volume, number of transactions, and number of items.

At least one embodiment includes tracking only a subset of volume types at a particular location .

The forecast structure comprises plural hierarchical levels of categories.

BRIEF DESCRIPTION OF THE DRAWINGS

- 25 The foregoing and other objects, features and advantages of the invention will be apparent from the following more particular description of preferred embodiments of the invention, as illustrated in the accompanying drawings in which like reference

009649122-082800

Fig. 2 is a schematic diagram which illustrates how the an embodiment of the present invention can allot time into one or more labor allocation classes.

Fig. 4 is a hierarchical chart providing an example in which three departments have been defined.

Fig. 6 is a schematic diagram illustrating an example of tasks and labor
15 standards defined within an embodiment of the present invention.

Fig. 8 is a table illustrating the relationship between the forecast structure and sources of volume data.

Fig. 10 is a block diagram of an embodiment of the forecaster.

Fig. 12 is a schematic diagram of an embodiment of the event ratio engine of

Fig. 13 is a schematic diagram of an embodiment of the volume forecasting of Fig. 10.

Fig. 14 is a view of a Forecast Event Ratio grid as used in an embodiment of the present invention.

Fig. 15 is a view of a Volume Ratio grid as used in an embodiment of the present invention.

5 Fig. 16 is a view of a Daily Ratio grid as used in an embodiment of the present invention.

Fig. 17 is a schematic diagram of an embodiment of the traffic pattern engine of Fig. 10.

Fig. 18 is a view of a Traffic Pattern Tab of Fig. 11.

Fig. 19 is a view of a dialog box which allows selection of specific dates for the Traffic Pattern Tab of Fig. 18.

Fig. 20 is a view of a dialog box which allows weighting of selection criteria for the traffic pattern engine of Fig. 17.

Fig. 21 is a schematic diagram of an embodiment of the workforce requirements
15 engine of Fig. 10.

Fig. 22 is a view of a dialog box which allows specification of user-defined forecast volumes.

Fig. 23 is a view of a workforce requirements window of the present invention.

Fig. 24 is a flowchart illustrating the major steps of the resource leveling
20 algorithm employed by an embodiment of the present invention.

Fig. 25 is a graph illustrating the concept of valleys, and valley depth and width, with an exemplary workforce requirement forecast.

DETAILED DESCRIPTION

A preferred embodiment of the present invention is a flexible workforce requirements forecaster. The flexibility of the forecaster is based on two partially dependent structures: a business structure, and a forecast structure. Each is defined by a user.

Business Structure

The "business structure" describes business operations. The business structure is made up of "labor allocation classes." A labor allocation class is a storage entity into which labor may be measured and tracked for periods of time.

5 In one embodiment, up to ten labor allocation classes can be defined.

Fig. 1 illustrates a sample business structure 10 which uses ten possible labor allocation classes 12, such as "Corporation", "Region", etc.

Labor allocation classes may be set up in a hierarchical format, for example, Business: Department: Job. Note, however, that in one embodiment, each class in the
10 business structure hierarchy may only have a single child class, i.e., multiple child classes are not allowed. For example, the Department class cannot have both Job and Machine classes under it.

Fig. 2 is a schematic diagram illustrating how time is allocated into one or more labor allocation classes. In this example, an employee John Smith has worked for eight
15 hours at the "machinist" job, on Work Order #12457, at the machine "Widget Maker."

The time 20 that John Smith worked, i.e., the eight hours, can be categorized into each one of these classes 22, i.e., Job, Work Order # and Machine, as indicated by the arrows 21. Labor allocated into a hierarchy is categorized and stored at the lowest level 22, and summed up in a consolidation of child values at each parent level.

20 Fig. 3 is a schematic diagram illustrating a "labor structure" 30 into which labor allocation classes are preferably organized, within the business structure 10. A "labor structure" is the collection of labor allocation classes, structured in hierarchical format, within which workforce requirements are generated and employees scheduled. A component of the labor structure is called a "labor level." For example, in the example
25 of Fig. 3, "Corporation" is the top labor level 12a in the labor structure 30.

In one embodiment, three specific labor levels are required for forecasting: the job level, the department level and the site level.

09649122-082800

SECRET

SECRET

SECRET

SECRET

SECRET

SECRET

SECRET

SECRET

Categories provide the logical structure of the user's forecast structure. They may exist as low-level units, or as consolidation points of lower-level categories. The forecast structure is always hierarchical. No low-level category may belong to more than a single higher-level category, one level directly above. The number of category levels may go as deep as desired.

Another user may wish to view only site transactions according to where they occurred, for example, at regular registers or express registers. This user might set up a forecast structure to have a parent category “Front End,” with children of “express” and “regular.”

Fig. 7 is a schematic diagram illustrating an exemplary forecast structure 50. Here, the root level node 52 corresponds to a particular site in the business structure.

The site "Burlington" has been divided into two categories, "Front End" 54a and "General Labor" 54b.

The "Front End" 54a, which represents the checkout area of the store, is further divided into "Express" 56a, representing express checkout lanes, and "Regular" 56b, representing all other checkout lanes. The "Express" 56a node is subdivided into "Cash" and "Credit", while the "Regular" 56b node is subdivided into "Cash", "Check", "Credit" and "Debit" 58.

Similarly, "General Labor" 54b is divided into "Bakery" 56c and "Deli" 56d. "Bakery" 56c is further subdivided into "Bread", "Cake" and "Pastry" categories 58 while "Deli" is subdivided into "Ham," "Roast Beef" and "Cheese" categories 58.

The "Front End" 54a, "Bakery" 56c and "Deli" 56d are category / department cross-reference points to the business structure. Note that the department assignment can exist at any level in the forecast structure. For example, in Fig. 7, "Front End" 54a is not at the same level in the forecast structure hierarchy as "Bakery" 56c and "Deli" 56d.

The makeup of the forecast structure depends on the level of detail that the user's volume collection systems can capture, as well as how the user wishes to "slice" the data and how drive labor. It is flexible enough to handle many different types of implementations.

20 The Forecast Structure and its Relation to Volumes

Once the forecast structure has been defined, a decision must be made as to which business volumes to track. Although between three and six different volumes are typically tracked, one embodiment of the present invention can track up to thirty-two different volumes. Examples of volumes in a retail environment are sales, items, customers, and weights.

Fig. 8 is a table 60 illustrating the relationship between the forecast structure 50a and sources of volume data for each category and for each type of volume. Categories

are portrayed vertically, while volume types 62, e.g., sales, items, customers, are listed along the horizontal axis. For each combination of volume and category in the forecast structure, the volume source is specified.

One type of volume data is direct feed, such as from a point of sale, on business volume for that category, e.g., Express.Cash Items. Another type of volume data is consolidated data, which may be summed from a category's children, or may be the basic level at which data is maintained. Alternatively, the user may specify, for a specific category, that no volume data is available, or the data may not be trustworthy.

Fig. 9 is a table 70 illustrating how labor standards can be defined based on those volumes. For example, for the job "Baker" in the "Bakery" department, the user might set up a standard called "Make Batter," which belongs to the "Bake Cake" task. The amount of labor required to fulfill this standard is driven by the number of cakes, i.e., cake items, that the site expects to sell.

Cross reference points between the business structure 36 and the forecast structure 50 can be specified, where the two structures are linked. These connections are made when a department from the business structure is assigned to a category in the forecast structure.

The forecaster sums workforce requirements from the job level up to the department level. Those workforce requirements can then be compared with the business volumes that they wish.

The "site" labor allocation class represents to top of the forecast structure.

A "department" is the connection point between the forecast structure and the labor structure. In the labor structure, users may attach a "department marker" to any labor allocation class above the job labor level. Jobs belong to their closest parent marked as a department. The forecaster uses the entities marked as departments as assignment points in the forecast structure for comparison of volume and labor data.

CALCULATION ENGINES

Fig. 10 is a block diagram of the forecaster of the present invention, which performs a variety of calculations. Each of these calculations is executed by a calculation engine. These calculation engines include a volume forecasting engine 80, a traffic pattern engine 100, an event ratio engine 120, and a workforce requirements engine 140.

The event ratio engine 120 reads an event calendar 125 in which "events" are defined, to produce event ratios 130. These event ratios 130, together with the historical volume data 89 which may be stored, for example, in a database, are input into the volume forecasting engine 80 and the traffic pattern engine 100.

The volume forecasting engine 80 provides a forecast 96 based on, for example, daily volume summaries. A forecast can be provided at every level of the forecast structure.

The traffic pattern engine 100 produces traffic pattern percentages 114 based on, for example, 15-minute periods, and using a separate calculation from the volume forecasting engine 80.

The volume forecast engine 80 and traffic pattern engine 100, although both volume forecasters, may use completely independent sets of data from the historical volume database 89. Alternatively, the data available to these two engines 80, 100 may partially or completely overlap. Thus, a user can select independent sets of historical data for the calculations.

The outputs 96, 114 of these engines are fed to the workforce requirements engine 140 which, given various business parameters 147, e.g., operating hours, and productivity standards 156, produces a workforce requirements forecast 158, which in turn can be fed to a scheduler.

Fig. 11 illustrates a sample screen presented to the user when the forecaster is invoked.

The default forecast date selection is one period after the current period. The system establishes the current date and looks up the next forecast period by looking at

009649122 082800

INDEX

INDEX

INDEX

INDEX

INDEX

INDEX

INDEX

INDEX

INDEX

5

10

15

the expected amount of business volume activity, and the actual amount that occurred.

20

25

5

10

15

are selected.

20

Forecast Event Ratio grid 180 which shows daily forecast event ratios for the week.

25

Volume Forecasting Engine

Fig. 13 is a schematic diagram of an embodiment of the volume forecasting engine 80 of Fig. 10. The volume forecasting engine 80 has the responsibility of calculating the prediction 96 of business volume for a future date. In one embodiment, one of two well-known forecasting models is used: "Exponential Smoothing" or "Daily Trend." The volume forecasting engine 80, using one of these two models, collects historical data points, making weighted adjustments for dates tagged as special "events" as discussed above, for each category in the forecast structure, and performs the appropriate calculations to predict the volume of business for a future day selected at 90. The user can then edit the volume forecast.

Inputs 88 to the volume forecasting engine 80 include volume history, such as sales, customers, transactions, etc., an event calendar, output from the event ratio engine, and a manager's plan which may be input manually.

The forecast structure 92 and a defined period 94, e.g., a 15-minute period, act as controls to the engine 80.

The resulting volume forecast is presented to the user in a volume forecast grid 176 (Fig. 11) of eight rows listing days of the forecast week plus volume total. One to thirty-two columns, depending on the configuration, list the volumes for which point of sales data is being collected. Depending on which node in the forecast category tree 172 has the focus, resulting forecast values may or may not exist in the volume columns based on what has been checked in Forecast Configurations.

The forecast values themselves are computed using either an Exponential Smoothing or Daily Trend calculation model. Which model to use can be specified for each category. The calculated value represents the aggregate business volume that the forecaster predicts for the selected future date based on a certain range of historical data. After calculation, the volume forecast is editable by the user. The volumes can be specified to have pre-configured lateral relationships with one another such that an edit to one value proportionally changes associated volume values.

In addition, edits to volume totals can uniformly change the daily values that contribute to them, so an edit to a category's volumes will change the total of the parent's volumes, as well as be distributed proportionally down to the child category volume forecasts.

5 There is an exception to this case, however. If both the parent volume and the children volumes are direct feeds, any edits to the parent volume's forecast will not be distributed down, nor will any edits to the child volume's forecasts consolidate up to the parent volume's forecast.

After editing a volume forecast cell, the user can choose to have the system
10 recalculate the labor forecast based on the updated value.

A volume forecast for a week can be adjusted during calculation by the application of forecast event ratios, as described above.

Within the Volume Forecast grid 176, this information is available to the user by right clicking a column heading.

15 “Volume ratios” are the relationships that the volumes have to one another. For instance, of interest to the user might be the ratio of sales to customer or the ratio of sales to items.

As illustrated in Fig. 15, after choosing to view Volume Ratios from a pop-up menu, the Volume Forecast grid is replaced with a Volume Ratio grid 185 which shows the ratios of all other volumes to the selected volume. In this example, CUSTOMERS is the selected column heading and the text appears when this option is selected. For the selected volume's column the number 1.00 will always be displayed.

The volume ratios can be edited by the user and the results are reflected in the corresponding cell when the view reverts back to the Volumes view 176. It is important to note, however, that edits to volume ratios do not have a lateral effect on other ratios the way edits to volume forecasts do, although others ratios that rely on the edited Volume ratio will naturally be recalculated.

[illegible][illegible][illegible]

067980

[illegible][illegible]

000006978

As Fig. 18 illustrates, although traffic patterns by day are calculated automatically when a forecast for a week is run, the Traffic Pattern Tab 230 allows the user to independently create, view, and manipulate the traffic patterns used by the forecaster to calculate workforce requirements.

5 Traffic patterns are presented in a graphical format driven by the forecast structure category selection made by the user. Two key graphs 232, 234 dominate the tab 230, with time represented along the horizontal axis, and the percentage of daily activity, e.g., sales, along the vertical axis.

In one embodiment, all applicable volume fields for a selected category are listed in box 236 using color codes. To view a particular volume field's traffic pattern, the user highlights the in box 236. A volume field, and a graph of the volume field's traffic pattern using the appropriate color is displayed. More than one volume can be selected by the user to display on the graphs.

The top graph 232 presents the traffic patterns for the selected categories and volume forecast, based on the checked dates in date grid 238. The forecast traffic pattern is a composite of the selected dates. Removing a checkmark from one of these dates changes the traffic pattern accordingly as that date will no longer contribute to the composite. This composite traffic pattern is used by the forecaster to calculate workforce requirements for any labor standards assigned to volumes within this pattern.

20 The bottom graph 234 displays the traffic pattern of a single date in the date grid 238. This allows a user to examine the difference between a single date's pattern and a composite pattern, which may or may not already include that date. Listed at the top of this graph 234 is a score 235 calculated for this historical date (November 8, 1997) based on user defined weighting criteria, discussed below.

25 Updated composite traffic pattern can be viewed by the user when the "Find [n]
best matches" button 240 is clicked. These updated patterns are committed to the
database via the "Save" button 242. Specific dates from the volume history database

can also be retrieved by the user and added to the composite traffic pattern via the "Find Specific Dates" button 241.

Fig. 19 illustrates a "Select Specific Dates" dialog box 250 which is displayed when the user selects the "Find Specific Date..." button 241 of Fig. 18. Dates selected from a calendar 252 are displayed in a list 254 with their corresponding scores 255.

Referring back to Fig. 18, the selection criteria "weight," i.e., the percent of influence a date has, used in scoring the historical traffic patterns, can be tuned by the user via the "Pattern Select Criteria" button 243.

Fig. 20 illustrates a "Traffic Pattern Selection Criteria" dialog box 260, which is displayed upon selection of the "Pattern Select Criteria" button 243 of Fig. 20.

Configurable settings include:

- Same Day of Week (Monday, Tuesday, etc.) 262,
- Nearest Day(Yesterday scores higher than day before) 264,
- Event Ratio(Days tagged with Event; e.g. Thanksgiving) 266, and
- Same Day Open/Close Time(08:00 am - 05:00pm) 268.

A "score" is calculated for a given date based on how closely it matches the criteria as weighted.

Workforce Requirements Engine

Fig. 21 is a schematic diagram of the workforce requirements engine 140 of Fig. 10. The fundamental task of the forecaster is to produce, for the scheduler, labor requirements based on the business that the site expects to receive, expressed in the volume forecast produced by either the volume forecast engine 80 or the traffic pattern engine 100. As discussed previously, one or more tasks are defined for each job. For each task, a number of labor standards are defined. These labor standards translate the amount of expected business activity for a site into actual labor requirements for each

job, expressed in either person-minutes, or people required to work during a period of time.

Labor standards come in two basic varieties. Labor standards can be fixed standards, i.e., standards that must be done regardless of business mix, such as overhead. Alternatively, labor standards can be variable, based on business activity, predicted either by the forecaster in the volume forecast, or by the user.

For each job, the forecaster calculates the amount of labor required for each labor standard by period of the day, and sums them up into a job total by period of the day. It then compares that total with the minimum and maximum number of employees that can work that job during that period of the day, and adjusts these totals if they do not meet the minimum and maximum requirements.

As with the volume forecast, a user may edit the labor forecast as desired. Users edit workforce requirements at the job level, by period of the day. When users edit workforce requirements, the forecaster does not reconcile the edits down to the task level. Instead, it stores the difference between the sum of the task-level requirements and the job level requirements as either unassigned or unallocated workforce requirements. These values are stored at the task level.

Inputs to the workforce requirements engine 140 include productivity standards, the output 96 of the volume forecast engine 80 and the output 114 of the traffic pattern engine 100. The user can specify which of the engines 80, 114 should drive each level of business activity

Controls include a user-defined forecast volume 150, workforce requirements adjustments 152, edits to the forecast and recalculated workforce requirements 154 and productivity standards 156.

25 The output comprises workforce requirements 158, by department, by job, by day and/or by time of day. This is the output of the forecaster which is fed to the scheduler.

06-07-2008

00000000000000000000

00000000000000000000

[illegible]

000000000000

[illegible]

00000000000000000000

00000000000000000000

[illegible]

00000000000000000000

[illegible][illegible]

00000000000000000000000000000000

00000000000000000000000000000000

[illegible][illegible]

080679

00000000000000000000000000000000

The dialog box 200 is initially driven by the specific forecast category node having focus on the forecast structure 172 (Fig. 11) and displays the jobs for the department attached to that node. The job control 224 allows the user to select the workforce requirements by job for viewing and editing after highlighting a specific day on the grid. After activating the dialog, however, the user may select a new department via the Department control 222. Both the Department and the Job controls have an “All” option that allows aggregate workforce requirements to be displayed in the grid, either for the entire department or for the entire enterprise.

15 A workforce driver is a component of every task definition, used in the
calculation of workforce requirements. The vast majority of workforce requirements are
based on fluctuating volume workforce drivers such as Sales, Items and Customers.
There may be times, however, when workforce requirements must be calculated for a
job for which no volume data is available. Examples include new employee training,
20 deliveries, inventory, and so on - jobs that are independent of the daily flow of business
volume.

Fig. 22 illustrates a dialog box 195 which is available from the Forecast Summary Tab to allow the user to specify values at runtime that can be used by the Workforce Requirements Engine 140 to calculate hours for non-volume-based jobs.

25 Task Level Consolidation

Schedule requirements are generated by consolidating. That is, the requirement for each standard for a task is first calculated, and then summed for the requirements for the task. This is then repeated at the job level.

Thus, the user is provided multiple levels of detail for calculation and thus has more flexibility in configuring productivity metrics. For example, a user might schedule employees by job, but prefer to assign tasks for each shift scheduled per job, while productivity standards are very detailed and granular. With only two levels, i.e., job and standard, the user has only the low level standard to assign. With three levels, i.e., job, task and standard, cashier shifts can be scheduled and tasks can be configured as groups of standards. For example, job: cashier; task: customer service; standard: ring item, tender cash, tender check, and so on.

The task level supports consolidation of standards into assignable groups.

Resource Leveling

There are some tasks that do not need to be performed at specific times. For example, floors may need to be swept just once in the morning and once in the afternoon, while the exact time is unimportant.

The workforce requirements calculation engine 140 can calculate the best time to perform the task. The present invention allows flexible tasks to be distributed within a certain window.

The problem which resource leveling addresses may best be seen as follows. Suppose the forecaster initially forecasts a requirement of 1.5 employees for some period. Obviously, the scheduler cannot schedule 1.5 people, therefore the requirement must be rounded to a whole number. However, if the requirement goes from, for example, 1.9 for a first period, to 0.9 for a second period, and finally to 2.9 for a third period, this is extremely difficult for a scheduler to handle.

Therefore, an embodiment of the forecaster identifies areas which are difficult for a scheduler. The forecaster fills gaps, i.e., valleys, by allocating resources to minimize rounding and simplify scheduling.

The resource leveling algorithm can also allocate time for breaks, which are
5 another kind of standard.

Fig. 24 is a flowchart illustrating the major steps of a resource leveling algorithm employed by an embodiment of the present invention.

At Step 401, a workforce requirements forecast is constructed, in person-minutes for the entire week (considering eight days in a week), i.e., for the 768 data points
10 corresponding to 768 15-minute periods, and for all non-resource-leveling tasks of some job.

For each resource-leveling task of the job, the loop comprising steps 403-423 is repeated.

It may not be desirable to calculate labor during certain periods, called “periods
15 of application” or POAs. Therefore, the user can define those periods of the week in which a standard may calculate requirements, and can turn the standard off for other portions of the day,

Thus, at Step 405, for each standard of the resource-leveling task, the week’s periods of application, i.e., those periods during which a resource-leveling task can be
20 assigned, are identified.

For each period, i.e., each 15-minute interval, of each POA as determined in Step 405, the loop comprising steps 407-421 is repeated.

At Step 408, all “remainders” are determined for each period. Remainders are calculated based on the following formula:

25
$$\text{Remainder} = (\text{Rounded Person Minutes} - \text{Forecasted Person Minutes})$$

09671R-08000

09671R-08000

09671R-08000

09671R-08000

[illegible]

09671R-08000

09671R-08000

09671R-08000

09671R-08000

09671R-08000

Fig. 25 is a graph which illustrates the valleys of an exemplary workforce requirement forecast. The vertical axis represents the number of employees required, while the horizontal axis represents time, partitioned into 15-minute intervals. The actual requirement values as forecasted before resource leveling are listed below the time axis, as well as represented in the graph as X's.

The peaks 450 of the graph have been identified and circled. Three valleys 452-456 have been identified.

The width and depth of each valley has been calculated and is shown in the figure. For example, Valley #1 452 has a width of 6, since it covers six 15-minute periods, from 10:00 am until 11:30 am. The depth is the difference between the values of the higher peak 450A defining the valley and the lowest point 451 in the valley, or $3.1 - 1.4 = 1.7$.

Referring again to Fig. 24, at Step 417, a rank is calculated for each valley as follows:

$$rank = (Depth / Width) * Cost$$

At Step 419, resource leveling is applied to fill the highest-ranked valley. That is, a resource leveling task is allotted as much as possible to time within the valley.

Dynamic Standard Assignment

In one embodiment of the present invention, different metrics, or standards, can be applied at different times, or to different locations, depending on conditions during the forecast period or at the different locations.

For example, suppose that supermarkets are classified as either "big" or "small." A newer, larger operation will likely generate the most volume, due for example, to economy of scale. Meanwhile, older and smaller supermarkets will have certain known inefficiencies. Each type of supermarket needs its own set of productivity metrics / standards.

It will be apparent to those of ordinary skill in the art that methods involved in the forecaster may be embodied in a computer program product that includes a computer usable medium. For example, such a computer usable medium can include a readable memory device, such as a hard drive device, a CD-ROM, a DVD-ROM, or a
5 computer diskette, having computer readable program code segments stored thereon. The computer readable medium can also include a communications or transmission medium, such as a bus or a communications link, either optical, wired, or wireless, having program code segments carried thereon as digital or analog data signals.

While this invention has been particularly shown and described with references
10 to preferred embodiments thereof, it will be understood by those skilled in the art that various changes in form and details may be made therein without departing from the scope of the invention encompassed by the appended claims.

008280-23754960